Programming Language Features for Refinement

Jason Koenig
Stanford University
jrkoenig@stanford.edu

K. Rustan M. Leino
Microsoft Research
leino@microsoft.com

Algorithmic and data refinement are well studied topics that provide a mathematically rigorous approach to gradually introducing details in the implementation of software. Program refinements are performed in the context of some programming language, but mainstream languages lack features for recording the sequence of refinement steps in the program text. To experiment with the combination of refinement, automated verification, and language design, refinement features have been added to the verification-aware programming language Dafny. This paper describes those features and reflects on some initial usage thereof.

0. Introduction

Two major problems faced by software engineers are the development of software and the maintenance of software. In addition to fixing bugs, maintenance involves adapting the software to new or previously underappreciated scenarios, for example, using new APIs, supporting new hardware, or improving the performance. Software version control systems track the history of software changes, but older versions typically do not play any significant role in understanding or evolving the software further. For example, when a simple but inefficient data structure is replaced by a more efficient one, the program edits are destructive. Consequently, understanding the new code may be significantly more difficult than understanding the initial version, because the source code will only show how the more complicated data structure is used.

The initial development of the software may proceed in a similar way, whereby a software engineer first implements the basic functionality and then extends it with additional functionality or more advanced behaviors. For example, the development of a library that provides binary decision diagrams (BDDs) may proceed as follows: the initial version may use a simple data structure; then, reductions are implemented; to facilitate quicker look-ups, hash-consing is added; caches are added to speed up commonly occurring operations; the garbage collection provided by the programming language is replaced by a custom allocator and collector that are specific to the needs of the BDD library; the functionality is extended to allow setting the variable ordering; watch dogs are added to monitor how the variable ordering is affecting performance; a system for automatically changing the variable ordering dynamically is added. Much before these steps have all been added, the software has reached considerable complexity and has become difficult to understand and costly to maintain.

If the design of a piece of software were explained from one software engineer to another, the explanations would surely be staged to explain the subsequent layers of complexity gradually. In this paper, we consider how a *programming language* can give software engineers the ability to write the source code in logical stages, in the same way that it may be explained in person. In particular, we describe our design of *refinement* features in the verification-aware programming language Dafny [16].

Stepwise program refinement, including data refinement, has been studied a great deal in the last several decades, see, e.g., [0, 3, 23, 26]. It provides a mathematical framework for gradually introducing complexity into a design. It has been implemented in software construction and modeling tools, like

Eerke Boiten, John Derrick & Steve Reeves (Eds.):
Refinement Workshop 2015 (Refine '15)
EPTCS 209, 2016, pp. 87–106, doi:10.4204/EPTCS.209.7

KIV [26], Atelier B, and Rodin [1]. However, the input to these tools take a larger departure from today's programming languages than we would like. From the language design perspective, we are looking for something more in the spirit of the Transform [8, 9] or SETL [27], but with tool support for both compilation and reasoning. As we report in this paper, we have found it difficult to design a usable set of features in the programming language. We are not ready to give up, however. Instead, we hope that our mixed experiences will inspire improved designs in the future.

We describe our design goals in Section 1. In Section 2, we describe Dafny's refinement features and illustrate these with small examples. Refining an instance of a simple class into an aggregate object that uses new instances of library-defined classes is difficult [6, 20]. In Section 3, we show how this is done in Dafny. In the last sections of the paper, we reflect on our experience, compare with related work, and conclude.

1. Design Goals

Our view is that the programming language is a software engineer's most important tool. Therefore, we think it is important to try to capture more of the design of a program into the program text itself. A program can use the constructs in a language to aid in making a design understandable, which is important both for development and maintenance.

A central pedagogical principle lies in presenting details at the right time, and this principle is manifested in many well-known programming facilities. Among these, *procedural abstraction*—whereby computational descriptions are divided into named, reusable routines—is perhaps the most universal. *(Interface and implementation) modules* provide another way to hide details, for example as the "one secret per module" guideline enunciated by Parnas [25]. In object-oriented software, *subclassing* gives a way to collect common behavior and to customize details in class-specific ways. In functional programming, *type parametricity* gives a way to operate over data without needing to be concerned with the specifics of the data, thus abstracting over the details. Cross-cutting details can also be introduced using *aspects*, which give a whole-program way to customize behavior [14]. We are hoping for refinement features that give yet another way to stage the complexity of a program. Whereas procedural abstraction allows layering of the call graph, refinement features aim to layer the logical complexity.

Our design goals are to provide:

- Programming in stages. We want the refinement features to allow logical, gradual introduction of details. We also want the result to be easier to understand than what alternative constructs provide today.
- A program structuring device. While it seems desirable for modules to provide strict information-hiding barriers, this is not usually a strong concern in procedural abstraction. Procedures internal to a module frequently factor out behavior without going as far as making sure callers and callees are entirely decoupled. For example, a change in the caller may require a change in the callee, and vice versa. It may be helpful to think of this as the *one-developer view*, where the one software developer is in control of both sides of the procedural abstraction boundary. Our aim is for the refinement features to have such a one-developer view. That is, we do not see a refinement boundary as a boundary that must support all sorts of uses. Instead, a refinement may be introduced as a program structuring device that just helps organize the program into logically staged pieces. A developer will not be shamed when making changes to the software that require changes to other sides of refinement boundaries. In particular, we will allow an initial design to *anticipate* further

refinements. For example, this makes it okay for the program to contain "shims" or "refinement points" that are to be filled in or referenced later.

- Program-like constructs. By considering refinement in the programming language, we are doing something that is different from mainstream programming languages. However, we do not want to stray too far—we want the result to still look, more or less, like mainstream programs today.
- Lightweight. We want the refinement features to be easy to use, without the need for bulky syntax that reduces understanding.
- Support reuse. Though we have the one-developer view, we do wish for constructs that lend themselves to reuse.
- Modular verification and compilation. We want it to be possible to reason about a piece of code without having to know the details of future refinements. Similarly, we want it to be possible to compile uses of an abstraction before all the details of the refinement have been decided.

2. Refinement in Dafny

Dafny is a programming language designed with reasoning in mind [16]. Its features include a repertoire of imperative and functional features. In addition, the language integrates constructs for specifying the intended behavior of programs, like pre- and postconditions, as well as features that facilitate stating lemmas and writing proofs. Dafny has a program verifier that checks that a program meets its given specifications. The integrated development environment (IDE) for the language constantly runs the verifier in the background in order to expedite feedback to the user [19].

Dafny's focus on reasoning and correctness makes it especially appealing as a testbed for introducing refinement features. We describe Dafny's refinement features in modules, in specifications of functions and methods, in method bodies, and across modules. In Section 3, we use a longer example to describe refinement features in classes.

2.0. Modules

A Dafny program is divided into *modules*. A module contains declarations of methods, functions, types (like inductive datatypes and instantiable classes, where classes themselves declare fields, methods, and functions), iterators, and nested modules. In addition, a module can *import* other modules.

One module can be declared to be a *refinement* of another module, as indicated by following the name of the new module with the keyword `refines` and the name of the module to be refined. The refining module is based on the module it refines, but it is a separate module. More precisely, the contents of the refined module is copied into the refining module, modulated by three kinds of *directives*:[0]

- **Extend** the refining module with additional declarations (for example, declare a new type or a new method)
- **Define** entities whose definition the refined module omitted (for example, define a previously opaque type or give a body of a previously body-less function)

[0] We speak about different *kinds* of directives only in order to explain the functionality provided in Dafny. The user never needs to name these directives when writing the program or running the verifier or compiler. Instead, which kind of directive to apply is implicit from the program text, as we shall see in examples.

```
abstract module A {
  type T
  function F(x: T): T
  function Twice(x: T): T
  { F(F(x)) }
}
module B refines A {
  type T = T'
  datatype T' = Leaf(int) | Node(T, T)
  function F...
  { match x
    case Leaf(w) => Leaf(w+1)
    case Node(left, right) => Node(F(left), F(right))
} }
```

Figure 0. Two example modules, one (`B`) declared as a refinement of the other (`A`).

- **Refine** previously given specifications (for example, strengthening postconditions) and previously given bodies of methods and (in one special case) functions

For example, Figure 0 shows a module `A` that declares an opaque type `T` and two functions, one of which (`F`) is body-less (that is, uninterpreted). In this example, `A` has been declared as *abstract*, which tells the compiler not to generate any code for `A`. Without the `abstract` keyword, the compiler would complain about the missing type definition and function body. (Note, a module that defines all its entities need not be abstract to be refined.) Module `B` is declared as a refinement of `A`. It extends `A` by declaring an inductive datatype `T'`. It also defines `T` to be a synonym for `T'` and it defines a body for `F`. Note that module `B` also contains function `Twice`, which is copied from module `A`. Also, recall that the presence of module `B` in the program does not affect module `A`; they are two separate modules.

In our example, we chose not to repeat the signature of `F`, but instead to use the syntax `...`. Dafny also allows the type signature of `F` to be repeated (allowing renamings of parameters) in the refining module.[1]

Because Dafny's refinement operates at the level of modules, it is possible to simultaneously refine a set of types. Compare this to the limited one-type refinements achievable by a disciplined use of subclassing in object-oriented languages.

One mechanical way to describe the refinement features in Dafny is to think of them as an elaborate template mechanism. However, Dafny restricts the use of the features to adhere to the standard *Principle of Semantic Refinement*, meaning that any client that is correct when using a module `A` is also guaranteed to be correct if `A` is replaced by any refinement of `A`. By analogy, object-oriented languages tend to provide a syntactic mechanism for subclassing, but do not insist that this mechanism be used only in accordance with *behavioral subtyping* [5, 21]. Since Dafny is equipped with a program verifier, its definition can afford to insist on following the Principle of Semantic Refinement (as opposed to just

[1] We have considered requiring the `...` syntax. This would always make it clear that the function is a refinement, and it would reduce the clutter and brittle nature of having to textually copy the signature. However, as even this simple example shows, the fact that the `...` syntax does not repeat the names of the parameters can also be confusing when looking at the body of the function ("What is `x`?").

providing a syntactic template mechanism).

Next, we will start to see how Dafny's restrictions preserve semantic refinements.

2.1. Specifications

Dafny distinguishes between *methods*, which are procedures with statements that can modify the program's heap, and *functions*, which are mathematical functions. Both can have specifications: pre- and postconditions (given by `requires` and `ensures` clauses), frame specifications (`modifies` clauses for methods and `reads` clauses for functions), and termination metrics (`decreases` clauses). A refinement module is allowed to add more `ensures` clauses, thus strengthening the postcondition of the method or function. In an analogous way, it would be sound to weaken preconditions and shrink frame specifications, but Dafny does not provide any syntax for doing so. Methods are allowed to be declared with `decreases *`, which says that the method is allowed to diverge. A refinement module is allowed to change this specification by giving a termination metric that proves termination.

For example, method `Max` in module `A` of Figure 1 has a weak specification. It allows the method to diverge, and if the method does terminate, the specification only says that the result (which is returned in the output parameter `m`) must not be smaller than the input parameters. Module `B` strengthens the postcondition of `Max` to say that the result is one of the input parameters. By giving a termination metric, it also says that `Max` terminates.

Because Dafny enforces the Principle of Semantic Refinement, the work of the verifier does not need to be repeated in refinement modules. In this example, when Dafny verifies module `A`, it checks that the implementation of `Max` meets the weak postcondition. When it verifies module `B`, it only checks that the implementation meets the additional postcondition and that the lexicographic tuple `x < y, x - y` strictly decreases with each recursive call.

From the specification of `Max` in `B`, Dafny also verifies the correctness of the assert statement in `Main`. Note, if no termination metric is given for `Max` in `B`, then it would inherit the "divergence allowed" from `A`; in that case, Dafny would complain that `Main`, which is not specified to allow divergence, is calling a possibly diverging method.

The two examples given so far show the directives Extend and Define. The Refine directive is more involved, as we describe next.

2.2. Statements

In what we have shown so far, a refining method can supply a body if the refined method omitted it. Dafny's Refine directive goes deeper than this and admits two kinds of change directives to a given method body:

- **Tighten Up** statements, to reduce nondeterminism
- **Superimpose** statements onto the refined method body, to introduce and modify additional program state

Since these directives apply to previously given statements or program points, there is a need to explain, as part of the program, where the directives are to apply. For this purpose, we have borrowed the *code skeletons* from Chalice [20]. Code skeletons work by listing in the refining method the changes from the refined method, when necessary mimicking the structure of the code in the refined method. We explain this functionality by example; see [20] for a full merge algorithm.

```
module A {
  method Max(x: int, y: int) returns (m: int)
    ensures x <= m && y <= m
    decreases *
  {
    if x == y {
      m := x;
    } else if x < y {
      m := Max(y, x);
    } else {
      m := Max(x-1, y);
      m := m + 1;
} } }
module B refines A {
  method Max...
    ensures m == y || m == x
    decreases x < y, x - y
  method Main() {
    var m := Max(10, 20);
    assert m == 20;
} }
```

Figure 1. A convoluted implementation for computing the maximum of two numbers. The specification of `Max` in module `B` strengthens the specification of `Max` in `A`.

Dafny offers several nondeterministic statements. These can be replaced by more deterministic statements. The replacement itself may incur some proof obligation, but previous proof obligations are not re-verified. For example, the "assign such that" statement `x :| P;` says to set variable `x` to any value satisfying the predicate `P` (there is a proof obligation that such an `x` exists) [18]. By the Tighten Up directive, this statement can be replaced by an ordinary assignment statement `x := E;`, incurring a proof obligation that `P` with `x` replaced by `E` holds.

For example, the pivot selection in QuickSort can first be implemented by a statement

```
var pivot :| lo <= pivot < hi;
```

and later refined to

```
var p0, p1, p2 := lo, (lo + hi) / 2, hi - 1;
if a[p2] < a[p0] {
  p0, p2 := p2, p0;
}
var pivot := if a[p1] < a[p0] then p0 else if a[p2] < a[p1] then p2 else p1;
```

This refinement superimposes statements that declare and assign to new local variables `p0`, `p1`, and `p2`, and then tightens up the assign-such-that statement to set `pivot` according to the "median of three" strategy. Dafny is able to distinguish the superimposition from the tighten up, since the merge algorithm matches the two assignments—one nondeterministic in the refined module and one deterministic in the refining module—to `pivot`. The refining module incurs a proof obligation that the value it assigns to `pivot` does indeed satisfy the condition indicated in the refined module.

The refining method is allowed to tighten up previous assignments and to modify superimposed state, but is not otherwise allowed to assign to previously declared variables. We refer to this as the New State Principle. For instance, the assignments to the new local variable `p0` in the example above are allowed and so is the assignment that tightens up the value of `pivot`, but `pivot` itself cannot be used as a temporary variable to hold any intermediate values.

Figure 2 shows another example where method `Abs` is specified to compute the absolute value of a given integer. Module `M0` uses a nondeterministic `if` statement that defines two control paths. One path sets the output parameter `a` to `x` and the other hopes to make `a` equal to `-x` using a loop. The method implementation establishes the postcondition only if the assumed conditions hold at the program points indicated. Note, for example, how the final assumption implies the last two conjuncts of the postcondition. Neither of the two `assume` statements is provable in module `M0`; not the first, because not enough information is known about `a` after the loop, and not the second, because the `if` statement allows control to flow through either branch.

Module `M1` in Figure 2 refines `M0` and tightens up the choice of which `if` branch to take. This allows the second `assume` statement to be turned into an `assert` statement. That is, the replacement of the `assume` with an `assert` incurs a proof obligation that the condition does hold at that program point, which is provable in module `M1`. The *elision statement*, `...;`, directs the merge algorithm to match any code sequence. Dafny implicitly inserts an elision statement at the end of every code block, that is, just before every "`}`", so all "`...;`" statements in the figure could have been omitted.

Dafny allows any number of refinement steps. The figure shows module `M1` being further refined by module `M2`. It turns the first `assume` statement into an `assert`, which is provable because of the added loop invariant. Note how expressions from the refined method are not repeated but instead replaced by "`...`".

```
abstract module M0 {
  method Abs(x: int) returns (a: int)
    ensures (a == x || a == -x) && x <= a && -x <= a
  {
    if * {
      a := x;
    } else {
      a := 0;
      var b := x;
      while b < 0 {
        a, b := a + 1, b + 1;
      }
      assume a == -x;
    }
    assume x <= a && -x <= a;
} }
abstract module M1 refines M0 {
  method Abs... {
    if 0 <= x {
      ...;
    } else {
      ...;
    }
    assert ...;
} }
module M2 refines M1 {
  method Abs... {
    if ... {
      ...;
    } else {
      ...;
      while ...
        invariant a + x == b <= 0
      { ...; }
      assert ...;
    }
    ...;
} }
```

Figure 2. An artificial example that shows several Tighten Up refinements. The `Abs` method in module `M0` postpones some proof obligations by introducing `assume` statements, and leaves some room for later deciding which `if` branch to take. Module `M1` tightens up the control flow and module `M2` fills in missing parts of the program's correctness argument.

Dafny provides a few statement refinement directives in addition to the ones we have shown by the example above. The general idea, as we have shown, is for the refining methods to mimic the structure of the method being refined, using `...;` to stand for elided code, superimposing new statements, and giving replacement statements that tighten up nondeterminism in the refined method. Dafny allows statements to be labeled (which outside of refinement is useful with `break` statements). Labels can be repeated in a refining method, which can occasionally be helpful as an aid for the merge algorithm.

With one exception, the refining method is not allowed to disrupt previous control flow. For example, the refining method is not allowed to add `break` statements that exit out a loop. The one exception is that new `return` statements are allowed. Dafny checks that the method's postcondition holds at those points in the refining method. This is useful, for example, if the refinement adds a cache or algorithmic support that enables a fast path in the method implementation.

Dafny includes two statements for the sole purpose of supporting refinements, the elision statement and the `modify` statement. The latter has the form

```
modify W { Body }
```

where `W` is a frame specification (which, like in a `modifies` clause, says which heap locations may be modified, and `{ Body }` is a block statement. Dafny treats the statement as the given block statement, but enforces that its heap modifications are in accordance with the frame specification. As we shall see in Section 3, the body of the `modify` statement can be postponed and defined in a refining method. If the body is omitted, the semantics of the statement is that of causing any arbitrary change permitted by the frame specification.

2.3. Clients

As one would expect from a language with a module system, Dafny allows a module to *import* other modules. This makes the declarations in the imported modules available to the importing module (the *client*) via qualified names. Since a module refinement gives rise to a separate module, an issue arises of how a client selects among the available refinements.

The basic import declaration has the form:

```
import X = M
```

where `M` is the name of a module defined elsewhere and `X` is a local name introduced as the qualifier when referring to declarations inside `M`. In the common case where one chooses a local name identical to the name of the imported module, the import declaration is abbreviated by just `import M`. The module import relation in a program must be acyclic. Moreover, an abstract module can be imported only by other abstract modules.

Consider a module `A0` and a refinement module `A1` (for brevity, we show the modules without contents here):

```
module A0 { }
module A1 refines A0 { }
```

A client module can choose to import either one of these by using `import A = A0` or `import A = A1`.

It is also possible to be less specific, by replacing the `=` with an `as`. The import declaration

```
import A as A0
```

says to use `A` as a local name for *some* module that *adheres* to `A0`, that is, whose contents (method bodies excluded) is a superset of the contents of `A0`. The eventual module imported can be `A0` itself, any

```
abstract module TotalOrder {
  type T
  predicate Below(x: T, y: T)
  lemma Transitive(x: T, y: T, z: T)
    requires Below(x, y) && Below(y, z)
    ensures Below(x, z)
  // other properties omitted from the figure
}
abstract module GenericSorting {
  import O as TotalOrder
  // sorting methods omitted from the figure
}
```

Figure 3. A sketch of a module that defines an ordering on a type `T`, and the import declaration of a module that makes use of that ordering.

refinement of `A0`, or in fact any other module that structurally is like `A0` or a refinement thereof.[2]

An "`as`" import in a module can be tightened up in a refinement module, as illustrated by the following example:

```
module B0 {
  import A as A0
}
module B1 refines B0 {
  import A = A1
}
```

Dafny checks that `A1` adheres to `A0`, which if `A1` is a module that refines `A0` is a trivial check. If `B1` wants to rely on the declarations in `A1`that were not in `A0`, but anticipates a further refinement of the imported module, then it can instead use an "`as`" import.

As an example, consider the modules in Figure 3. Module `TotalOrder` defines a type `T`, a relation `Below` on that type, and an unproved lemma (that is, an axiom) that states a property of `Below`. We have omitted lemma declarations for other properties that might also be useful. Module `GenericSorting` imports some module like `TotalOrder`. This lets it define methods (omitted in the figure) that sort values of type `O.T` according to the order `O.Below`.

Figure 4 shows refinements of the modules in Figure 3. In particular, module `IntOrder` defines `T` to be a synonym for `int`, defines `Below` to be the less-or-equal ordering on integers, and gives a (trivial) proof that the property `Transitive` holds. Module `IntSorting` refines `GenericSorting` by tightening up the import declaration. Consequently, the refining module will contain copies of the refined module's methods, but specialized for integers.

Note that the features we discuss in this paper do not give rise to dynamic dispatch (like the *traits* feature in Dafny does [2]). There is no relation between refinement modules that can be exploited dynamically at run time.

[2] It is also possible to combine the `as` and `=` imports: the declaration `import A as A0 default A1` is essentially treated like `import A as A0` by the verifier and as `import A = A1` by the compiler.

```
module IntOrder refines TotalOrder {
  type T = int
  predicate Below... { x <= y }
  lemma Transitive... { }
  // proofs of other properties omitted from the figure
}
module IntSorting refines GenericSorting {
  import O = IntOrder
}
```

Figure 4. The modules of Figure 3 specialized to integers.

3. Classes and Data Refinement

An important part of giving a simple description of a program lies in choosing variables with simple types. For example, sets and maps are often used, but details of how to represent such sets and maps are not. The systematic coordinate transformation from such abstract data structures to more efficient ones is called *data refinement* (among many other sources, see [3, 8, 9]). Getting data refinement to work in the presence of classes is difficult, because of encapsulation issues with references to dynamically allocated objects [6, 20].

To present a small example that gives brief taste of the essential problem, consider the following class:

```
class Interval { var width: int }
```

With appropriate refinement rules, it is known how such a data structure can be refined into, say:

```
class IntervalEndPoints {
  var start: int
  var end: int
}
```

where `width` is represented as the difference `end - start`. In this case, the fields `start` and `end` are introduced in the refinement, and thus by the New State Principle, these assignments to `start` and `end` are allowed in the refinement

As an alternative refinement that involves reuse of library components, suppose a library contains a class `Cell`:

```
class Cell { var data: int }
```

We may now consider a refinement like this:

```
class IntervalCell {
  var start: Cell
  var end: Cell
}
```

where `width` is represented as `end.data - start.data`. However, the soundness of this kind of refinement is much more involved. First, although the fields `start` and `end` are introduced in the refining

class, the field `data` was available already in the program being refined, and thus the simple New State Principle does not apply. Instead, allowing the refinement to modify the values of `start.data` and `end.data` requires more elaborate refinement rules. The intuition is that the particular objects referenced by `start` and `end` were never allocated in the program being refined, so `start.data` and `end.data` in effect constitute new state. For more information about this problem, along with solutions, see [6, 20].

Dafny uses idioms of *dynamic frames* to specify behavior of the heap [13, 15, 28]. The basic idea is to programmatically keep track of the set of individual objects that as an aggregate provide the behavior of the abstract object. This *representation set* is often stored in a field

```
ghost var Repr: set⟨object⟩
```

The field is declared as *ghost*, meaning it is used only for reasoning about the program. The compiler erases ghost code, so at run time they appear only in spirit [7, 16].

Dafny does not have any specific data refinement or *transform* constructs [8, 9], but the combination of ghost code, superimposition, and a directive that allows predicates to be strengthened gives the ability to introduce data structures in stages. We proceed by giving an example, introduced in several stages.

3.0. A Counter Specification

In the first stage, we give a specification of a very simple class, see Figure 5. Abstractly, the class represents a counter, whose value is stored in ghost field `N`. The class also declares a field `Repr` as described above and a predicate `Valid()` that holds when the object is in its steady state. That is, the body of `Valid()` (omitted in module `M0`) is the *class invariant* of `Counter` [22]. (We explain the keyword `protected` in Section 3.2.)

The class also declares a constructor and two methods. The last postcondition of each of these is the familiar specification. The other parts of the specifications are exactly the idiomatic Dafny dynamic-frame specifications for a constructor, a mutating method, and a query method, respectively.[3] The occurrences of `Valid()` express that the class invariant holds on all method boundaries. The conjuncts that mention `fresh` say that any objects that the constructor or mutating method add to the representation set are freshly allocated, which is important for callers to know [15]. Finally, the `modifies` clauses say that the constructor is only allowed to modify the state of the object being constructed (which for the purpose of these specifications is treated as if it was allocated immediately before the constructor is called) and that `Inc` is allowed to modify the state of any object in the set `Repr`. In addition, every constructor and method is allowed to allocate new object and modify their state.

Module `M0` gives a client's view of the `Counter` class. The refinements that follow give the implementation of the class.

3.1. Defining Bodies

We now define the predicate, constructor, and methods by giving them bodies, see Figure 6. By separating modules `M0` and `M1`, we simply achieve what in a language like, say, Modula-3 would be done by writing a module interface and a module implementation [24].

Predicate `Valid()` says that the receiver is always part of the representation set, and the `null` reference is not. The constructor needs to add to `Repr` all objects that are to be part of the object's initial representation. The details of this set are determined in further refinements. The constructor body in `M1`

[3] By marking a class with the `{:autocontracts}` attribute, a pre-pass of the Dafny verifier will fill in the idiomatic parts of specifications automatically, thus reducing clutter in the program text.

```
abstract module M0 {
  class Counter {
    ghost var N: int
    ghost var Repr: set⟨object⟩
    protected predicate Valid()
      reads this, Repr
    constructor ()
      modifies this
      ensures Valid() && fresh(Repr - {this})
      ensures N == 0
    method Inc()
      requires Valid()
      modifies Repr
      ensures Valid() && fresh(Repr - old(Repr))
      ensures N == old(N) + 1
    method Get() returns (n: int)
      requires Valid()
      ensures n == N
} }
```

Figure 5. A module that gives the standard, idiomatic dynamic-frames specification of a simple class.

```
abstract module M1 refines M0 {
  class Counter {
    protected predicate Valid... {
      this in Repr && null !in Repr
    }
    constructor ... {
      ghost var repr: set⟨object⟩ :| null !in repr && fresh(repr);
      N, Repr := 0, repr + {this};
    }
    method Inc... {
      N := N + 1;
      modify Repr - {this};
    }
    method Get... {
      n :| assume n == N;
} } }
```

Figure 6. A refinement of the module in Figure 5, containing a simple `Counter` class. Module `M1` defines the bodies in terms of the ghost fields `Repr` and `N`.

anticipates these further additions by introducing a local variable `repr`, which it allows to contain any set of newly allocated objects.

Similarly, method `Inc` uses the `modify` statement, anticipating that further refinements will want to do state changes of any representation object other than `this`. (Note that by the New State Principle, a refinement can still modify fields of `this`, provided those fields are declared in the refinement module.)

Method `Get` sets output parameter `n` to `N`, but in a somewhat roundabout way. First, in order to allow refinements to change how `n` is computed, `Get` uses an assign-such-that statement rather than a more straightforward assignment statement `n := N;`. Second, since `n` is not a ghost variable, the right-hand side of the assignment ordinarily must not depend on ghost variables like `N`. Use of the keyword `assume` in the assign-such-that statement indicates to Dafny that this statement is not intended to be compiled, so Dafny relaxes the ordinary restriction on ghost dependencies.[4]

3.2. An Implementation

We introduce a concrete implementation of the counter. We assume there is some `Library` module with a `Cell` class and use two instances of this class. The value of the counter, `N`, is represented as the difference between the `data` field of these two objects, see Figure 7.

The class is extended with the declaration of new fields `c` and `d`. By superimpositions, the constructor straightforwardly allocates two `Cell` objects and assigns these to `c` and `d`. The constructor then tightens up the value assigned to `repr`.

Method `Inc` defines a body for the `modify` statement and method `Get` tightens up the assignment to `n` by assigning it a value computed from non-ghost fields. To discharge the proof obligation that the `modify` body modifies only what is allowed by the frame specification and the proof obligation incurred by the tighten-up directive, it is necessary to have a stronger class invariant. In particular, the former proof obligation requires `c in Repr` and the latter requires `N == c.data - d.data`. In addition, the well-formedness checks for the statements introduced require `c` and `d` to be non-`null`.

Strengthening the class invariant comes down to changing the definition of `Valid()` to a stronger predicate. This is dicey, because `Valid()` appears in preconditions and it is not sound to strengthen preconditions in general. Inside the refining module, the verifier can arrange to re-verify proof obligations that involve establishing `Valid()` or assuming `!Valid()`. But what about client modules that were verified against the module being refined? Such verifications would also have to be redone, which means verification would no longer be modular. For this reason, Dafny allows a predicate to be strengthened only if it is marked as `protected`, which means the predicate's exact definition will never be revealed outside the module. Consequently, other modules cannot rely on the exact definition of the predicate, and so they are insensitive to any changes of it.

The syntax for this Predicate Strengthening directive is the same as that to Define a predicate. In other words, if a refining module gives a body for a predicate that already had a body, the effect is that of changing the definition of the predicate to the conjunction of the two bodies. This is allowed only for predicates marked as `protected`.

It is possible to continue refining `M2` into a subsequent module with more state, but doing so requires changes to `M2` that let it anticipate further refinements. For example, module `M2` may need to introduce another variable like `repr` in the constructor and to superimpose another `modify` statement in method `Inc`—in the same way that `M1` anticipated the further refinements given by `M2`. Methodologically, the fact

[4] The fact that `assume` has the desired effect here is rather coincidental. It would probably be better to change Dafny to allow ghost variables in right-hand sides of assign-such-that statements in abstract modules.

```
module M2 refines M1 {
  import Library
  class Counter {
    var c: Library.Cell
    var d: Library.Cell
    protected predicate Valid... {
      c in Repr && d in Repr &&
      c ≠ d &&
      N == c.data - d.data
    }
    constructor ... {
      c := new Library.Cell(0);
      d := new Library.Cell(0);
      ghost var repr: set⟨object⟩ := {c,d};
    }
    method Inc... {
      ...;
      modify ... {
        c.data := c.data + 1;
      }
    }
    method Get... {
      n := c.data - d.data;
} } }
```

Figure 7. A further refinement of the module that defines the `Counter` class. This refinement implements the counter in terms of two dynamically allocated `Cell` objects.

```
module M3 refines M2 {
  class Counter {
    method Get... {
      if d.data == 0 { return c.data; }
      ...;
} } }
module M4 refines M3 {
  class Counter {
    protected predicate Valid... { d.data == 0 }
} }
```

Figure 8. Module `M3` adds a fast path to the `Get` method of module `M2`, and module `M4` strengthens predicate `Valid()` to demonstrate that `d.data == 0` is in fact an invariant of the class.

that further refinements may require changes to the module to be refined is justified (and even considered normal) by the one-developer view. An analogous situation arises in object-oriented programming, when a new subclass needs an existing class to introduce dynamically dispatched calls to a new method.

Rather than taking our example in the direction of adding more state, we will in the next subsection illustrate the gist of a performance optimization that does not require further data refinements.

3.3. A Performance Optimization

Figure 8 shows a module `M3` that refines `M2`. It applies a directive only to method `Get`, into whose body it superimposes an `if` statement. The new code sets up a fast path in the event that the stated condition holds. In Dafny, the `return` statement with argument expressions has the effect of assigning the expressions to the output parameters and then returning from the method. Since output parameters do not fall under the New State Principle, the refinement is normally not allowed new assignments to them; however, as this is a useful and harmless case, the implicit assignments to output parameters that happens as part of a superimposed `return` statement are allowed.

In our simple example, the fast path we introduced will not give rise to any actual performance improvement unless the compiler realizes that `d.data == 0` actually always holds (in which case the condition does not need to be tested in the emitted code). To illustrate how refinements could help give that information to the compiler, we can strengthen the class invariant further, see module `M4` in Figure 8.

It is worth mentioning once more that `M2`, `M3`, and `M4` are three separate modules. Dafny checks the refinement among these successive modules, but does not relate the classes they define. In particular, classes `M2.Counter`, `M3.Counter`, and `M4.Counter` are three separate types and are not subclasses of one another.

This completes our illustration of how a class can be built in stages. Looking back at Figures 5 through 8, the elisions are such that the refinements from module to module stand out. A user can inspect what any ellipsis stands for by placing the mouse pointer above the ellipsis in the Dafny IDE, upon which the elided information will be displayed as a hover text.

4. Experience, Evaluation, and Related Work

We have used the refinement features in Dafny for a number of toy programs. Although the provided directives can accomplish the usual refinement tasks, our impression is that refinement works more smoothly on paper than in our language. Things that, due to hand waving, may be simple to achieve on paper (like the problem solved by the local variable `repr` in Figure 6) look more clumsy in our language design.

One could argue that useful formal-methods techniques are also useful if applied informally, that is, without actually carrying through the proofs. This argument leads to asking if our superimposition and tighten-up directives are useful devices for program structuring. Here, too, it is not clear that our design gets a good score. For one, the fact that one needs to declare another module in order to stage some refinements can feel bulky.

A similar bulkiness issue also arises in Event-B [0] implemented in the Rodin tool [1], where all unchanged events have to be copied into the file that contains the subsequent refinement. An alternative is given by the "refinement layer" annotations in Civil [12]. These allow several stages of refinement to be given in a single source text. The verifier processes a given program once for each declared layer, suitably ignoring the declarations of higher-numbered layers.

When authoring or reading a sequence of refinements, one sometimes wants to see only the changes from one module to the next and sometimes wants to see the full resulting program. Our elision statements only address the former, and our IDE's hover text does not adequately address the latter. We had chosen the elision statements under a rather traditional view that a program is a printable piece of program text. A more modern or even futuristic view would be to let the sequence of refinements appear as layered text in the IDE. A user could then be given various ways to input and read the program. The refinement tools KIV [26] and Rodin [1] have embraced the idea that the IDE can manage the program better than a line-by-line editor can. We hope such environments will also be developed for languages that look more similar to today's mainstream languages than KIV and Event-B do.

A desirable scenario to support in staged program development is to write a program in an abstract way and then replace the operations on certain variables with other, more efficient operations on alternative variables. This is a central goal of the *transform* by Gries et al. [8, 9]. At first, the rather syntactic match-and-replace rules in these transforms appear brittle. But given that this is a scenario we would like to support smoothly, and given that we are buying into the one-developer-view idea of anticipating refinements, we would be interested in incorporating the transform into Dafny.

The Dafny design that a refinement module creates a separate module is a feature in some cases. For example, it allows multiple refinements of the `TotalOrder` module in Figure 3, each one of which can benefit from reuse. It has also been used to define common processing of different services in the IronFleet project, which was authored in Dafny [11]. But we have also seen it make the common interface-implementation pattern rather verbose, since it requires a refinement module when a client wants to tighten up which implementation gets used for the abstract module it `as`-imported. (We have started exploring an alternative module design wherein every abstract module has a default refinement module.)

An early tool for machine-assisted program development lets the user apply *refinement tactics* to massage a formal specification into code [10]. The tactics applied are recorded and can be displayed. Moreover, the IDE allows a user to expand a sub-specification to see what it has been refined into; conversely, the details can be elided to instead show just the more abstract sub-specification that they implement. We would wish for an IDE that keeps track of the program-derivation tree in this way. However, we also note that Dafny provides greater flexibility in introducing correlated transformations (like

the addition of an `assume` statement in each branch of an `if` to justify some other refinement transformation after the `if` statement), and it is not clear how these can be presented with an equally simple IDE.

Despite many shortcomings in our language design, the current refinement features in Dafny have been useful in some complex examples. One such example is a break-down of the Schorr-Waite algorithm into stages. More precisely, the proof obligations, loop invariants, and ghost variables used in the proof were broken down into a sequence of refinements that seems to separate concerns in a desired way. Another example is the formalization of the Cloudmake algorithm [4]. It introduces some axiomatized functions and later uses refinements to prove the feasibility of those axioms. Interestingly enough, these examples use the refinement features mostly to structure *proofs*, not to structure the executable statements of the program.

The examples in our paper can be tried online at `http://rise4fun.com/Dafny/{4FH, 74s9, jrJQ, jX5Y, n07}`. Additional examples can be found in the Dafny test suite at `http://dafny.codeplex.com`. A video of a SPLASH 2012 keynote with live demos is also available online [17].

5. Concluding Remarks

We have described the refinement features in version 1.9.5 of Dafny. While far from perfect, we have combined refinement and automated verification into a programming language. We hope that use of our system will inspire further exploration and innovation in incorporating refinement features in day-to-day programming languages.

Acknowledgments

We are grateful to Lindsay Groves and Mark Utting for feedback at the REFINE 2015 workshop, and to the referees for their detailed readings and constructive feedback.

References

[0] Jean-Raymond Abrial (2010): *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, doi:10.1017/CBO9781139195881.

[1] Jean-Raymond Abrial, Michael Butler, Stefan Hallerstede, Thai Son Hoang, Farhad Mehta & Laurent Voisin (2010): *Rodin: An Open Toolset for Modelling and Reasoning in Event-B*. *International Journal on Software Tools for Technology Transfer*, doi:10.1007/s10009-010-0145-y.

[2] Reza Ahmadi, K. Rustan M. Leino & Jyrki Nummenmaa (2015): *Automatic Verification of Dafny Programs with Traits*. In Rosemary Monahan, editor: *Formal Techniques for Java-like Programs, FTfJP 2015*, ACM, doi:10.1145/2786536.2786542.

[3] Ralph-Johan Back & Joakim von Wright (1998): *Refinement Calculus: A Systematic Introduction*. Graduate Texts in Computer Science, Springer-Verlag, doi:10.1007/978-1-4612-1674-2.

[4] Maria Christakis, K. Rustan M. Leino & Wolfram Schulte (2014): *Formalizing and Verifying a Modern Build Language*. In Cliff B. Jones, Pekka Pihlajasaari & Jun Sun, editors: *FM 2014: Formal Methods — 19th International Symposium*, *Lecture Notes in Computer Science* 8442, Springer, pp. 643–657, doi:10.1007/978-3-319-06410-9_43.

[5] Krishna Kishore Dhara & Gary T. Leavens (1996): *Forcing Behavioral Subtyping through Specification Inheritance*. In H. Dieter Rombach, T. S. E. Maibaum & Marvin V. Zelkowitz, editors: *18th International Conference on Software Engineering*, IEEE Computer Society, pp. 258–267, doi:10.1109/ICSE.1996.493421.

[6] Ivana Filipović, Peter O'Hearn, Noah Torp-Smith & Hongseok Yang (2010): *Blaming the client: on data refinement in the presence of pointers*. *Formal Aspects of Computing* 22(5), pp. 547–583, doi:10.1007/s00165-009-0125-8.

[7] Jean-Christophe Filliâtre, Léon Gondelman & Andrei Paskevich (2014): *The Spirit of Ghost Code*. In Armin Biere & Roderick Bloem, editors: *Computer Aided Verification — 26th International Conference, CAV 2014*, *Lecture Notes in Computer Science* 8559, Springer, pp. 1–16, doi:10.1007/978-3-319-08867-9_1.

[8] David Gries & Jan Prins (1985): *A New Notion of Encapsulation*. In: *Proceedings of the ACM SIGPLAN 85 Symposium on Language Issues in Programming Environments*, *SIGPLAN Notices 20* 7, ACM, pp. 131–139, doi:10.1145/800225.806834.

[9] David Gries & Dennis Volpano (1990): *The Transform — a New Language Construct*. *Structured Programming* 11(1), pp. 1–10.

[10] Lindsay Groves, Raymond Nickson & Mark Utting (1992): *A Tactic Driven Refinement Tool*. Technical Report CS-TR-92/5, Victoria University of Wellington, Department of Computer Science, doi:10.1007/978-1-4471-3550-0_14.

[11] Chris Hawblitzel, Jon Howell, Manos Kapritsos, Jacob R. Lorch, Bryan Parno, Michael L. Roberts, Srinath T. V. Setty & Brian Zill (2015): *IronFleet: Proving Practical Distributed Systems Correct*. In Ethan L. Miller & Steven Hand, editors: *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP 2015*, ACM, pp. 1–17, doi:10.1145/2815400.2815428.

[12] Chris Hawblitzel, Erez Petrank, Shaz Qadeer & Serdar Tasiran (2015): *Automated and Modular Refinement Reasoning for Concurrent Programs*. In Daniel Kroening & Corina S. Pasareanu, editors: *Computer Aided Verification — 27th International Conference, CAV 2015 (Part II)*, *Lecture Notes in Computer Science* 9207, Springer, pp. 449–465, doi:10.1007/978-3-319-21668-3_26.

[13] Ioannis T. Kassios (2006): *Dynamic Frames: Support for Framing, Dependencies and Sharing Without Restrictions*. In Jayadev Misra, Tobias Nipkow & Emil Sekerinski, editors: *FM 2006: Formal Methods, 14th International Symposium on Formal Methods*, *Lecture Notes in Computer Science* 4085, Springer, pp. 268–283, doi:10.1007/11813040_19.

[14] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Videira Lopes, Jean-Marc Loingtier & John Irwin (1997): *Aspect-Oriented Programming*. In Mehmet Aksit & Satoshi Matsuoka, editors: *ECOOP'97 — Object-Oriented Programming, 11th European Conference*, *Lecture Notes in Computer Science* 1241, Springer, pp. 220–242, doi:10.1007/BFb0053381.

[15] K. Rustan M. Leino (2009): *Specification and verification of object-oriented software*. In Manfred Broy, Wassiou Sitou & Tony Hoare, editors: *Engineering Methods and Tools for Software Safety and Security*, *NATO Science for Peace and Security Series D: Information and Communication Security* 22, IOS Press, pp. 231–266. Summer School Marktoberdorf 2008 lecture notes.

[16] K. Rustan M. Leino (2010): *Dafny: An Automatic Program Verifier for Functional Correctness*. In Edmund M. Clarke & Andrei Voronkov, editors: *LPAR-16*, *Lecture Notes in Computer Science* 6355, Springer, pp. 348–370, doi:10.1007/978-3-642-17511-4_20.

[17] K. Rustan M. Leino (2012): *Staged Program Development*. SPLASH 2012 keynote, InfoQ video, http://www.infoq.com/presentations/Staged-Program-Development, doi:10.1145/2384716.2384719.

[18] K. Rustan M. Leino (2015): *Compiling Hilbert's ε Operator*. In Ansgar Fehnker, Annabelle McIver, Geoff Sutcliffe & Andrei Voronkov, editors: *LPAR-20. 20th International Conferences on Logic for Programming, Artificial Intelligence and Reasoning — Short Presentations*, *EPiC Series in Computer Science* 35, EasyChair, pp. 106–118.

[19] K. Rustan M. Leino & Valentin Wüstholz (2014): *The Dafny Integrated Development Environment*. In Catherine Dubois, Dimitra Giannakopoulou & Dominique Méry, editors: *Proceedings 1st Workshop on Formal Integrated Development Environment, F-IDE 2014*, *EPTCS* 149, pp. 3–15, doi:10.4204/EPTCS.149.2.

[20] K. Rustan M. Leino & Kuat Yessenov (2012): *Stepwise refinement of heap-manipulating code in Chalice*. *Formal Aspects of Computing* 24(4–6), pp. 519–535, doi:10.1007/s00165-012-0254-3.

[21] Barbara Liskov & Jeannette M. Wing (1994): *A Behavioral Notion of Subtyping*. *ACM Transactions on Programming Languages and Systems* 16(6), doi:10.1145/197320.197383.

[22] Bertrand Meyer (1988): *Object-oriented Software Construction*. Series in Computer Science, Prentice-Hall International.

[23] Carroll Morgan (1990): *Programming from Specifications*. Series in Computer Science, Prentice-Hall International.

[24] Greg Nelson, editor (1991): *Systems Programming with Modula-3*. Series in Innovative Technology, Prentice-Hall, Englewood Cliffs, NJ.

[25] D. L. Parnas (1972): *On the criteria to be used in decomposing systems into modules*. *Communications of the ACM* 15(12), pp. 1053–1058, doi:10.1145/361598.361623. Reprinted as `www.acm.org/classics/may96/`.

[26] Wolfgang Reif (1992): *The KIV System: Systematic Construction of Verified Software*. In Deepak Kapur, editor: *Automated Deduction — CADE-11, 11th International Conference on Automated Deduction*, *Lecture Notes in Computer Science* 607, Springer, pp. 753–757, doi:10.1007/3-540-55602-8_218.

[27] J. T. Schwartz, R. B. K. Dewar, E. Dubinsky & E. Schonberg (1986): *Programming with Sets: An Introduction to SETL*. Texts and Monographs in Computer Science, Springer, doi:10.1007/978-1-4613-9575-1.

[28] Jan Smans, Bart Jacobs, Frank Piessens & Wolfram Schulte (2008): *Automatic Verifier for Java-Like Programs Based on Dynamic Frames*. In José Luiz Fiadeiro & Paola Inverardi, editors: *Fundamental Approaches to Software Engineering, 11th International Conference, FASE 2008*, *Lecture Notes in Computer Science* 4961, Springer, pp. 261–275, doi:10.1007/978-3-540-78743-3_19.